

A Note on Quantum Security for Post-Quantum Cryptography

Fang Song
Department of Combinatorics & Optimization
and Institute for Quantum Computing
University of Waterloo

Abstract

Shor’s quantum factoring algorithm and a few other efficient quantum algorithms break many classical crypto-systems. In response, people proposed post-quantum cryptography based on computational problems that are believed hard even for quantum computers. However, security of these schemes against *quantum* attacks is elusive. This is because existing security analysis (almost) only deals with classical attackers and arguing security in the presence of quantum adversaries is challenging due to unique quantum features such as no-cloning.

This work proposes a general framework to study which classical security proofs can be restored in the quantum setting. Basically, we split a security proof into (a sequence of) classical security reductions, and investigate what security reductions are “quantum-friendly”. We characterize sufficient conditions such that a classical reduction can be “lifted” to the quantum setting.

We then apply our lifting theorems to post-quantum signature schemes. We are able to show that the classical generic construction of hash-tree based signatures from one-way functions and a more efficient variant proposed in [BDH11] carry over to the quantum setting. Namely, assuming existence of (classical) one-way functions that are resistant to efficient quantum inversion algorithms, there exists a quantum-secure signature scheme. We note that the scheme in [BDH11] is a promising (post-quantum) candidate to be implemented in practice and our result further justifies it. Actually, to obtain these results, we formalize a simple criteria, which is motivated by many classical proofs in the literature and is straightforward to check. This makes our lifting theorem easier to apply, and it should be useful elsewhere to prove quantum security of proposed post-quantum cryptographic schemes. Finally we demonstrate the generality of our framework by showing that several existing works (Full-Domain hash in the quantum random-oracle model [Zha12b] and the simple hybrid arguments framework in [HSS11]) can be reformulated under our unified framework.

1 Introduction

Advances in quantum information processing and quantum computing have brought about fundamental challenges to cryptography. Many classical cryptographic constructions are based on computational problems that are assumed hard for efficient classical algorithms. However, some of these problems, such as factoring, discrete-logarithm and Pell’s equation, can be solved efficiently on a quantum computer [Sho97, Hal07]. As a result, a host of crypto-systems, e.g, the RSA encryption scheme that is deployed widely over the Internet, are broken by a quantum attacker.

A natural countermeasure is to use *quantum-resistant* assumptions instead. Namely, one can switch to other computational problems which appear hard to solve even on quantum computers, and construct cryptographic schemes based on them. Examples include problems in discrete lattices [MR09, Pei09] and hard coding problems [Sen11]. We can also make generic assumptions such as the existence of one-way functions that no efficient quantum algorithms can invert. This leads to the active research area termed

post-quantum cryptography [BBD09]. Nonetheless, quantum-resistant assumptions alone do not immediately imply quantum security of a scheme, due to other fundamental issues that could be subtle and easily overlooked.

First of all, we sometimes fail to take into account possible attacks unique to a quantum adversary. In other words, classical definition of security may not capture the right notion of security in the presence of quantum attackers¹. For example, many signature schemes are designed in the random-oracle (RO) model, where all users, including the attacker, can query a truly random function. This is meant to capture an idealized version of a hash function, but in practice everyone instantiate it by him/herself with a concrete hash function. As a result, when we consider quantum attacks on these schemes, there seems no reason not to allow a quantum adversary to query the random-oracle in quantum superposition. This leads to the so called quantum random-oracle model [BDF⁺11], in which we need to reconsider security definitions (as well as the analysis consequently) [Zha12b, Zha12a, BZ13].

A more subtle issue concerns security proofs, which may completely fall through in the presence of quantum attacks. Roughly speaking, one needs to construct a *reduction* showing that if an efficient attacker can successfully violate the security requirements of a scheme then there exists an efficient algorithm that breaks some computational assumption. However, a classical reduction may no longer work (or make sense at all) against quantum adversaries. A key classical technique, which encounters fundamental difficulty in the presence of quantum attackers, is called *rewinding*. Loosely speaking, rewinding arguments consist of a mental experiment in which an adversary for a scheme is executed multiple times using careful variations on its input. This usually allows us to gain useful information in order to break the computational assumption. As first observed by van de Graaf [vdG97], rewinding seems impossible with a quantum adversary since running it multiple times might modify the entanglement between its internal state and an outside reference system, thus changing the system's overall behavior. This issue is most evident in cryptographic protocols for zero-knowledge proofs and general secure computation. There has been progress in recent years that develops quantum rewinding techniques in some special cases [Wat09, Unr12], and a few classical protocols are proven quantum-secure [DL09, LN11, HSS11]. Hallgren et al. [HSS11] also formalized a family of classical security proofs against efficient adversaries that can be made go through against efficient quantum adversaries under reasonable computational assumptions. Despite these efforts, however, still not much is known in general about how to make classical security proofs go through against quantum attackers.

This note revisits these issues for post-quantum cryptography based on computational assumptions, focusing on simple *primitives* such as signatures, encryptions and identifications, where constructions and analysis are usually not too complicated (compared to secure computation protocols for example). In this setting, the issues we have discussed seem less devastating. For instance, rewinding arguments appear only occasionally, for example in some lattice-based identification schemes [Lyu08, Lyu09]. Usually rewinding is not needed for the security proof. Nonetheless, it is still crucial to pinning down proper security definitions against quantum attacks, as illustrated in the quantum random-oracle example above. In addition, just because there are no rewinding arguments, does not mean that we can take for granted that the security reduction automatically holds against quantum attackers. Very often in the literature of post-quantum cryptography, a construction based on some quantum-resistant assumption is given together with a security proof for *classical* attackers only. The construction is then claimed to be quantum-secure without any further justification. In our opinion, this is not satisfying and quantum security of these schemes deserves a more careful treatment.

CONTRIBUTIONS. The main contribution of this note is a general framework to study which classical security proofs can be restored in the quantum setting. A security proof can be split into (a sequence of) classical security reductions, and we investigate what reductions are “quantum-friendly”. Recall that

¹Although our focus is security against computationally bounded attackers, this issue is also relevant in the *statistical* setting. There are classical schemes, which are proven secure against unbounded classical attackers, broken by attackers using quantum entanglement [CSST11].

informally a reduction transforms an adversary of one kind to another. We distinguish two cases, *game-preserving* and *game-updating* reductions.

A game-preserving reduction is one such that the transformation still makes sense, i.e., syntactically well-defined, for quantum adversaries. In this case we propose the notion of *class-respectful* reductions which ensures in addition that the adversary obtained from the transformation indeed works (e.g., it is an efficient quantum algorithm and successfully solves some problem). Motivated by the structure of security reductions that occur in many post-quantum cryptographic schemes, we further characterize a simple criteria, which is straightforward to check. This makes the lifting theorem easier to apply, and should be useful to prove quantum security for many other schemes not restricted to the applications we show later in this note.

On the other hand, a game-updating reduction captures the case that the classical reduction no longer makes sense, as illustrated by the quantum random-oracle model. This is usually more difficult to analyze. We propose *translatable* reductions, which essentially reduces the problem to the game-preserving case. The basic idea is to introduce an “interpreter”, so that the classical reduction becomes applicable to a quantum adversary with the translation by the interpreter. In both cases, we show in our lifting theorems that a reduction can be constructed in the quantum setting if there is a classical reduction that is respectful or translatable respectively.

We apply our framework to prove quantum security of some hash-based signature schemes. Specifically, we show that the classical generic construction of hash-tree based signature schemes from one-way functions carries over to the quantum setting, assuming the underlying one-way function is quantum-resistant. This is also true for a more efficient variant proposed in [BDH11] assuming quantum-resistant pseudorandom functions, which in turn can be based on quantum-resistant one-way functions from known results. This scheme is a promising (post-quantum) candidate to be implemented in practice and our result further justifies it. Moreover, we give an alternative proof for the security of a general construction of signatures based on trapdoor permutations called Full-Domain hash in the quantum random-oracle model. We also show that an existing framework in the context of cryptographic protocols that characterizes a class of “quantum-friendly” classical security proofs (simple hybrid arguments [HSS11]) fits our framework. These demonstrate the generality of our framework.

REMARKS. Our framework (e.g., definitions of games and reductions) should look natural to people familiar with the provable-security paradigm. It should also be straightforward (or even obvious for experts) to verify the characterizations of “quantum-friendly” reductions in our lifting theorems. The purpose of this note, however, is to at least make people become more serious and cautious and to encourage further research, in addition to suggesting one possible formal framework to reason about the security of post-quantum cryptography against quantum attacks. Likewise, it may be just a tedious exercise to work through the classical proof for hash-based signatures and convince oneself it is indeed quantum-secure. Nonetheless, this can be done in a more abstract and rigorous way using our framework. We hope that our framework can be applied elsewhere to analyze quantum security of other post-quantum cryptographic constructions. Ideally, in some easy cases, it would serve as a tool to automate the routine parts, so that whoever designs a new scheme should be able to make some simple checks and conclude its quantum security.

OTHER RELATED WORKS. There are a few works that study systematically what classical proofs or statements can be “lifted” to the quantum setting in the context of multi-party secure computation. Unruh in [Unr10] showed that any classical protocol that is secure in the statistical setting, i.e., against computationally *unbounded* adversaries, under a strong *universal-composable* notion is also statistically secure in an analogous quantum universal-composable model. Fehr et al. [FKS⁺13] considered *reducibility* between two-party cryptographic tasks in the quantum setting. For example, one can ask if there is a secure protocol for oblivious transfer assuming two parties can perform bit commitments securely. They showed that in most cases, the reducibility landscape remains unchanged in the quantum setting under the very same classical protocols. However, there are cases that classical reducibility no longer holds quantumly, and sometimes new relations can be established using quantum protocols.

The formalization of games, reductions and other terms in this note is influenced by a lot of classical literatures on game-playing proofs [GM84, Yao82, KR01, BR06, Sho05, Hal05]. Recent developments, especially the framework of code-based game-playing proofs [Hal05, BR06] have motivated automated tools for proving security [Bla08, BGZB09, Stu09, BGHB11]. Our treatment of computational assumptions is also inspired by the line of works classifying complexity-theoretic intractability assumptions [Nao03, HH09, Pas11, GW11].

2 Preliminary

BASIC NOTATIONS. For $m \in \mathbb{N}$, $[m]$ denotes the set $\{1, \dots, m\}$. We use $n \in \mathbb{N}$ to denote a *security parameter*. The security parameter, represented in unary, is an implicit input to all cryptographic algorithms; we omit it when it is clear from the context. Quantities derived from protocols or algorithms (probabilities, running times, etc) should be thought of as functions of n , unless otherwise specified. A function $f(n)$ is said to be negligible if $f = o(n^{-c})$ for any constant c , and $\text{negl}(n)$ is used to denote an unspecified function that is negligible in n . We also use $\text{poly}(n)$ to denote an unspecified function $f(n) = O(n^c)$ for some constant c . When D is a probability distribution, the notation $x \leftarrow D$ indicates that x is a sample drawn according to D . When D is a finite set, we implicitly associate with it the uniform distribution over the set. If $D(\cdot)$ is a probabilistic algorithm, $D(y)$ denotes the distribution over the output of D corresponding to input y . We will sometimes use the same symbol for a random variable and for its probability distribution when the meaning is clear from the context. Let $\mathbf{X} = \{X_n\}_{n \in \mathbb{N}}$ and $\mathbf{Y} = \{Y_n\}_{n \in \mathbb{N}}$ be two ensembles of binary random variables. We call \mathbf{X}, \mathbf{Y} *indistinguishable*, denoted $\mathbf{X} \approx \mathbf{Y}$, if $|\Pr(X_n = 1) - \Pr(Y_n = 1)| \leq \text{negl}(n)$.

MACHINE MODELS. We model classical parties as interactive Turing machines, which are probabilistic polynomial-time (PPT) by default. Quantum machines are modelled following that of [HSS11]. A *quantum interactive machine* (QIM) M is an ensemble of interactive circuits $\{M_n\}_{n \in \mathbb{N}}$. For each value n of the security parameter, M_n consists of a sequence of circuits $\{M_n^{(i)}\}_{i=1, \dots, \ell(n)}$, where $M_n^{(i)}$ defines the operation of M in one round i and $\ell(n)$ is the number of rounds for which M_n operates (we assume for simplicity that $\ell(n)$ depends only on n). We omit the scripts when they are clear from the context or are not essential for the discussion. M (or rather each of the circuits that it comprises) operates on three registers: a state register \mathbf{S} used for input and workspace; an output register \mathbf{O} ; and a network register \mathbf{N} for communicating with other machines. The size (or running time) $t(n)$ of M_n is the sum of the sizes of the circuits $M_n^{(i)}$. We say a machine is polynomial time if $t(n) = \text{poly}(n)$ and there is a deterministic classical Turing machine that computes the description of $M_n^{(i)}$ in polynomial time on input $(1^n, 1^i)$. When two QIMs M and M' interact, their network register \mathbf{N} is shared. The circuits $M_n^{(i)}$ and $M'_n^{(i)}$ are executed alternately for $i = 1, 2, \dots, \ell(n)$. When three or more machines interact, the machines may share different parts of their network registers (for example, a private channel consists of a register shared between only two machines; a broadcast channel is a register shared by all machines). The order in which machines are activated may be either specified in advance (as in a synchronous network) or adversarially controlled.

3 Defining Games and Reductions

This section introduces a formal definition of reductions, which captures the type of security reductions that we care mostly about. It builds upon a basic notion of games.

We use *game* G to denote a general probabilistic process between two players: the challenger \mathcal{C} initiates the interaction with the other player, call it an adversary \mathcal{A} . After several rounds of communication, \mathcal{C} outputs one bit `succ` or `fail` indicating success or failure of the game. We define the game value of G with an adversary \mathcal{A} to be the probability that \mathcal{C} outputs `succ`, and denote it $\omega_G(\mathcal{A})$. Typically in a game G , \mathcal{C}

is efficient, i.e., a poly-time classical or quantum machine. Very often we want to analyze the game when the adversary is restricted to a class of machines \mathcal{C} (e.g., poly-time classical machines). We write $G(\mathcal{C})$ to indicate this case, and define $\omega_G(\mathcal{C}) := \max\{\omega_G(\mathcal{A}) : \mathcal{A} \in \mathcal{C}\}$. Sometimes we denote \hat{G} to stress a game defined for quantum machines. We describe below as an example the standard forgery game of existentially unforgeable signatures under (adaptive) chosen-message-attacks (EU-CMA) [GMR88, KL07].

Existential-Forgery Game G^{FOR}

Signature scheme: $\Pi = (\text{KGen}, \text{Sign}, \text{Vrfy})$.

- \mathcal{C} generates $(pk, sk) \leftarrow \text{KGen}(1^n)$. Send pk to adversary \mathcal{A} .
- \mathcal{A} can query signatures on messages $\{m_i\}$. \mathcal{C} returns $\sigma_i := \text{Sign}(sk, m_i)$. These messages can be chosen adaptively by \mathcal{A} .
- \mathcal{A} outputs (m^*, σ^*) . If $\text{Vrfy}(pk, (\sigma^*, m^*)) = 1$ and $m^* \notin \{m_i\}$, \mathcal{C} outputs succ. Otherwise output fail.

There are many variants of this game which will be used later in this note. For example, we denote the game in which \mathcal{A} is allowed to query at most one signature $G^{\text{OT-FOR}}$. $G^{\text{RO-FOR}}$ denotes the game where a random-oracle is available to both parties, and if the random-oracle can be accessed in quantum superposition we denote the game $G^{\text{QRO-FOR}}$.

We define a reduction \mathcal{R} as a 3-tuple $(G^{\text{ext}}, \mathcal{T}, G^{\text{int}})$. There are an external (explicit) game G^{ext} and an internal (implicit) game G^{int} , and an additional party \mathcal{T} called the *transformer*. Loosely speaking, \mathcal{T} transforms an adversary \mathcal{A} in G^{int} to an adversary in G^{ext} . Specifically, \mathcal{T} takes an adversary's machine \mathcal{A} as input and outputs the description of an adversary in G^{ext} . We distinguish *black-box* and *non-black-box* reductions, with a focus on black-box reductions. In a black-box reduction, \mathcal{A} is provided as a black-box, which means that the transformation does not look into the codes and inner workings of the adversary. Whereas in a non-black-box reduction, \mathcal{R} has the explicit description of \mathcal{A} . We denote $\mathcal{T}(\mathcal{A})$ as the resulting adversary in G^{ext} that is “transformed” from \mathcal{A} by \mathcal{T} . In the black-box setting, the output of \mathcal{T} will always be of the form $T^{\mathcal{A}}$, i.e., an oracle machine with access to \mathcal{A} . Note that T is the same for all \mathcal{A} , and it emulates an execution of G^{int} with \mathcal{A} . However, in general T needs not to run the game as in a real interaction. For instance, it can stop in the middle of the game and start over (i.e., rewind).

PROPERTIES OF A REDUCTION. To make a reduction meaningful, we describe below a few properties that we may want a reduction to hold. Let \mathfrak{A} and \mathfrak{B} be two classes of machines.

- **\mathfrak{A} -compatible** reductions. We say \mathcal{R} is \mathfrak{A} -compatible, if $\forall \mathcal{A} \in \mathfrak{A}$, $G^{\text{int}}(\mathcal{A})$ and $G^{\text{ext}}(\mathcal{T}(\mathcal{A}))$ are well defined. Namely \mathfrak{A} and $\mathcal{T}(\mathcal{A})$ respect the specifications of the games.
- **$(\mathfrak{A}, \mathfrak{B})$ -consistent** reductions. We say \mathcal{R} is $(\mathfrak{A}, \mathfrak{B})$ -consistent, if \mathcal{R} is \mathfrak{A} -compatible and $\forall \mathcal{A} \in \mathfrak{A}$, $\mathcal{T}(\mathcal{A}) \in \mathfrak{B}$. When we write a reduction as $(G^{\text{ext}}(\mathfrak{B}), \mathcal{T}, G^{\text{int}}(\mathfrak{A}))$ or $\mathcal{R}(\mathfrak{A}, \mathfrak{B})$ in short, the reduction is assumed to be $(\mathfrak{A}, \mathfrak{B})$ -consistent. Note that if \mathcal{R} is black-box, it must hold that $T^{\mathfrak{A}} \subseteq \mathfrak{B}$.
- **Value-dominating**. We say \mathcal{R} is value-dominating if $\omega_{G^{\text{ext}}}(\mathcal{T}(\mathcal{A})) = \omega_{G^{\text{ext}}}(\mathcal{T}(\mathcal{B}))$ whenever $\omega_{G^{\text{int}}}(\mathcal{A}) = \omega_{G^{\text{int}}}(\mathcal{B})$.
- **$(\alpha_{\text{succ}}, \mathfrak{A})$ -effective** reductions. Let $\alpha_{\text{succ}} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be some function. We say \mathcal{R} is α_{succ} -effective on \mathfrak{A} if $\omega_{G^{\text{ext}}}(\mathcal{T}(\mathcal{A})) \geq \alpha_{\text{succ}}(\omega_{G^{\text{int}}}(\mathcal{A}))$. If this holds for any $\mathcal{A} \in \mathfrak{A}$, we call \mathcal{R} $(\alpha_{\text{succ}}, \mathfrak{A})$ -effective.
- **$(\alpha_{\text{time}}, \mathfrak{A})$ -efficient** reductions. Let $\alpha_{\text{time}} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be some function. We say \mathcal{R} is α_{time} -efficient if $\text{TIME}(\mathcal{T}(\mathcal{A})) \leq \alpha_{\text{time}}(\text{TIME}(\mathcal{A}))$ for any $\mathcal{A} \in \mathfrak{A}$.

Effective and efficient reductions are often used in combination, especially when we are concerned with tightness of a reduction. In that case, α_{succ} and α_{time} may depend on both $\text{TIME}(\mathcal{A})$ and $\omega_{G^{\text{int}}}(\mathcal{A})$. This paper will focus on effectiveness only. We often abuse notation and use α_{succ} as a scalar if this causes no confusion. We stress that these properties talk about the output machine of \mathcal{T} on \mathcal{A} (e.g., $\mathcal{T}(\mathcal{A})$ lies in a specific class, or it runs in time comparable to that of \mathcal{A}), however we do not restrict the computational power of \mathcal{T} , though it is typically efficient. The reason is that for our purpose, we only need to show existence of an adversary for G^{ext} with nice properties.

4 Quantum-Friendly Security Reductions: A General Framework

In this section, we attempt to propose a general framework to study which classical proofs still hold when the adversaries become quantum. Consider a classical cryptographic scheme Π . To analyze its security against efficient classical attacks (in the provable-security paradigm), one typically proceeds as follows:

1. Formalizing some security requirement by a game G^{int} . Typically we are concerned about security against a particular class of attackers (e.g., PPT machines), so we restrict the game G^{int} to a class \mathfrak{A} . We also associate a value $\varepsilon_{\mathfrak{A}} \in (0, 1]$ with the game, which upper bounds the success probability that any adversary in \mathfrak{A} wins the game. Namely we require that $\omega_{G^{\text{int}}}(\mathfrak{A}) \leq \varepsilon_{\mathfrak{A}}$. We denote this security requirement as $(G^{\text{int}}(\mathfrak{A}), \varepsilon_{\mathfrak{A}})^2$.
2. Formalizing some computational assumption by another game G^{ext} . Similarly the assumption is assumed to hold against a specific class of machines, so we restrict the game to a class \mathfrak{B} , and require that $\omega_{G^{\text{ext}}}(\mathfrak{B}) \leq \varepsilon_{\mathfrak{B}} \in (0, 1]$. Denote the computational assumption as $(G^{\text{ext}}(\mathfrak{B}), \varepsilon_{\mathfrak{B}})$.
3. Constructing an $(\mathfrak{A}, \mathfrak{B})$ -consistent reduction $\mathcal{R} = (G^{\text{ext}}(\mathfrak{B}), \mathcal{T}, G^{\text{int}}(\mathfrak{A}))$. Security follows if the reduction is in addition α_{succ} -effective with $\alpha_{\text{succ}} \geq \varepsilon_{\mathfrak{B}}/\varepsilon_{\mathfrak{A}}$. This implies if there exists an $\mathcal{A} \in \mathfrak{A}$ with $\omega_{G^{\text{int}}}(\mathcal{A}) > \varepsilon_{\mathfrak{A}}$ (i.e., \mathcal{A} breaks the security requirement), there is an adversary $\mathcal{T}(\mathcal{A}) \in \mathfrak{B}$ such that $\omega_{G^{\text{ext}}}(\mathcal{T}(\mathcal{A})) \geq \alpha_{\text{succ}} \cdot \omega_{G^{\text{int}}}(\mathcal{A}) > \varepsilon_{\mathfrak{B}}$ (i.e., it breaks the computational assumption).

Now we want to know if the classical security reductions are “quantum-friendly” so that we can claim that the scheme is secure against quantum attacks. We need to reconsider each step of the classical analysis in the quantum setting (See Table 1 for a comparison between classical provable-security and quantum provable-security for a scheme.). Let $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ be two classes of quantum machines. We adapt G^{int} and define $(\hat{G}^{\text{int}}(\hat{\mathfrak{A}}), \varepsilon_{\hat{\mathfrak{A}}})$. It is supposed to capture some security requirement against quantum attackers in $\hat{\mathfrak{A}}$, and we require that $\omega_{\hat{G}^{\text{int}}}(\hat{\mathfrak{A}}) \leq \varepsilon_{\hat{\mathfrak{A}}}$. Likewise, we adapt G^{ext} to a game \hat{G}^{ext} , which should formalize a reasonable computational assumption $(\hat{G}^{\text{ext}}(\hat{\mathfrak{B}}), \varepsilon_{\hat{\mathfrak{B}}})$ against quantum adversaries. Then we can ask the fundamental question (still informal):

Can we “lift” \mathcal{R} to the quantum setting?

Namely, is there a reduction $\hat{\mathcal{R}}(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ that preserves similar properties as $\mathcal{R}(\mathfrak{A}, \mathfrak{B})$?

To answer this question, we distinguish two cases. In the simpler case, \hat{G} are syntactically identical to G . Namely, $\hat{G}^{\text{ext}}(\hat{\mathfrak{B}})$ (resp. $\hat{G}^{\text{int}}(\hat{\mathfrak{A}})$) is just G^{ext} (resp. G^{int}) restricted to the quantum class $\hat{\mathfrak{B}}$ (resp. $\hat{\mathfrak{A}}$). In particular, this means that G^{ext} and G^{int} are still the right games that capture a computational assumption and some security requirement. We call this case *game-preserving*. In contrast, as illustrated by the quantum random-oracle example, \hat{G} may change and this leads to a more complicated case to analyze. We call it *game-updating*. In the following subsections, we investigate in each case what reductions can be lifted to the quantum setting, and hence are quantum-friendly.

²Sometime we write $(G^{\text{int}}(\mathfrak{A}), \varepsilon_{\mathfrak{A}})_{\Pi}$ to emphasize the specific scheme we are dealing with, though it is usually clear from the context.

Table 1: Components of classical and quantum provable-security for a classical construction.

	Classical Provable-Security	Quantum Provable-Security
Security Requirement	$(G^{\text{int}}(\mathfrak{A}), \varepsilon_{\mathfrak{A}})$	$(\hat{G}^{\text{int}}(\hat{\mathfrak{A}}), \varepsilon_{\hat{\mathfrak{A}}})$
Computational Assumption	$(G^{\text{ext}}(\mathfrak{B}), \varepsilon_{\mathfrak{B}})$	$(\hat{G}^{\text{ext}}(\hat{\mathfrak{B}}), \varepsilon_{\hat{\mathfrak{B}}})$
Reduction	$\mathcal{R}(\mathfrak{A}, \mathfrak{B})$	$\xrightarrow{?} \hat{\mathcal{R}}(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$

4.1 Lifting Game-Preserving Reductions

Let $\mathcal{R}(\mathfrak{A}, \mathfrak{B}) = (G^{\text{ext}}(\mathfrak{B}), \mathcal{T}, G^{\text{int}}(\mathfrak{A}))$ be a classical reduction. Let $\hat{G}^{\text{ext}}(\hat{\mathfrak{B}})$ and $\hat{G}^{\text{int}}(\hat{\mathfrak{A}})$ be extended games in the quantum setting that are restricted to classes of quantum machines $\hat{\mathfrak{B}}$ and $\hat{\mathfrak{A}}$. We consider the case that \hat{G} and G are the same in this section. We want to know if there is a reduction $\hat{\mathcal{R}}(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ that preserves nice properties of \mathcal{R} . Since we are dealing with the same games applied to different classes of machines, one may expect that simple tweaks on \mathcal{R} should work. This intuition is indeed true to some extent, which we formalize next.

Definition 4.1 (*G-equivalent machines*). Two machines M and N are called *G-equivalent* if $\omega_G(M) \equiv \omega_G(N)$.

Definition 4.2 (*[G, C]-realizable classical machines*). A classical machine M is called *[G, C]-realizable*, if there is a machine $N \in \mathcal{C}$ s.t. $\omega_G(M) = \omega_G(N)$. We denote $E_G(\mathcal{C})$ as the collection of classical machines that are *[G, C]-realizable*.

We put forward *class-respectful* reductions as a template for quantum-friendly reductions in the game-reserving case.

Definition 4.3 (β - $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -respectful reductions). Let \mathcal{R} be a classical reduction $(G^{\text{ext}}(\mathfrak{B}), \mathcal{T}, G^{\text{int}}(\mathfrak{A}))$. We say \mathcal{R} is β - $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -respectful for some $\beta \in \mathbb{R}^+$ if the following hold:

1. $(\beta, \hat{\mathfrak{A}})$ -**extendable**: \mathcal{R} is $E_{G^{\text{int}}}(\hat{\mathfrak{A}})$ -compatible and $(\beta, E_{G^{\text{int}}}(\hat{\mathfrak{A}}))$ -effective. That is, $\forall \mathcal{A} \in E_{G^{\text{int}}}(\hat{\mathfrak{A}})$, $G^{\text{ext}}(\mathcal{T}(\mathcal{A}))$ and $G^{\text{int}}(\mathcal{A})$ are well-defined³, and $\omega_{G^{\text{ext}}}(\mathcal{T}(\mathcal{A})) \geq \beta(\omega_{G^{\text{int}}}(\mathcal{A}))$.
2. $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -**closed**: \mathcal{R} is $(E_{G^{\text{int}}}(\hat{\mathfrak{A}}), E_{G^{\text{ext}}}(\hat{\mathfrak{B}}))$ -consistent. Namely, $\forall \mathcal{A} \in E_{G^{\text{int}}}(\hat{\mathfrak{A}})$, $\mathcal{T}(\mathcal{A}) \in E_{G^{\text{ext}}}(\hat{\mathfrak{B}})$.

The theorem below follows (almost) immediately from this definition.

Theorem 4.4 (Quantum lifting for game-preserving reductions). *If $\mathcal{R}(\mathfrak{A}, \mathfrak{B})$ is β - $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -respectful, then there exists an $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -consistent reduction $\hat{\mathcal{R}}(\hat{\mathfrak{A}}, \hat{\mathfrak{B}}) := (G^{\text{ext}}(\hat{\mathfrak{B}}), \hat{\mathcal{T}}, G^{\text{int}}(\hat{\mathfrak{A}}))$ that is $(\beta, \hat{\mathfrak{A}})$ -effective.*

Proof. Consider any $\hat{A} \in \hat{\mathfrak{A}}$. Let \mathcal{A} be a classical machine such that \mathcal{A} is G^{ext} -equivalent to \hat{A} . Since \mathcal{R} is $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -closed, we know that $\mathcal{T}(\mathcal{A}) \in E_{G^{\text{ext}}}(\hat{\mathfrak{B}})$ and hence there is a machine $N_{\hat{A}} \in \hat{\mathfrak{B}}$ s.t. $\omega_{G^{\text{int}}}(N_{\hat{A}}) = \omega_{G^{\text{int}}}(\mathcal{T}(\mathcal{A}))$. Define $\hat{\mathcal{T}}$ to be a quantum machine such that, given $\hat{A} \in \hat{\mathfrak{A}}$, outputs $N_{\hat{A}}$. Namely $\hat{\mathcal{T}}(\hat{A}) := N_{\hat{A}}$. Let $\hat{\mathcal{R}} := (G^{\text{ext}}(\hat{\mathfrak{B}}), \hat{\mathcal{T}}, G^{\text{int}}(\hat{\mathfrak{A}}))$. Clearly $\hat{\mathcal{R}}$ is $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -consistent due to the way we defined $\hat{\mathcal{T}}$. It is also $(\beta, \hat{\mathfrak{A}})$ -effective because $\omega_{G^{\text{ext}}}(\hat{\mathcal{T}}(\hat{A})) = \omega_{G^{\text{ext}}}(\mathcal{T}(\mathcal{A})) \geq \beta(\omega_{G^{\text{int}}}(\mathcal{A})) = \beta(\omega_{G^{\text{int}}}(\hat{A}))$. \square

³Most classical games we deal with are actually well-defined for all machines. But we explicitly state this requirement in case of some artificial examples.

To apply the theorem, we need to check the two conditions of respectful reductions. The “extendability” condition is usually easy to verify. However, the “closure” property can be challenging and subtle, depending on the classes of players we care about. We will be mostly interested in poly-time machines. Namely let $\mathfrak{A} = \mathfrak{B}$ be poly-time classical machines and $\hat{\mathfrak{A}} = \hat{\mathfrak{B}}$ be the collection of poly-time quantum machines, denote it by \mathcal{Q} . In this case, we propose a simple criteria that is easy to check in existing classical security reductions. When combined with a few other easily verifiable conditions, we can show class-respectful reductions. This in a way justifies a common belief that most post-quantum schemes are indeed quantum-secure, due to some simple form in their classical security reductions which seem “quantum-friendly”.

Let $\mathcal{R} = (G^{\text{ext}}, \mathcal{T}, G^{\text{int}})$ be a classical black-box reduction. We say that \mathcal{R} is *straight-line* if the output machine of \mathcal{T} on \mathcal{A} , which as before is denoted $T^{\mathcal{A}}$, runs \mathcal{A} in straight-line till completion. Namely, other than the flexibility of choosing \mathcal{A} 's random tape, T behaves exactly like a honest challenger in G^{int} when it invokes \mathcal{A} . This type of reduction, due to its simple structure, is amenable to getting lifted.

Theorem 4.5 (Straight-line reduction: a useful condition for class-closure). *Let $\mathcal{R} = (G^{\text{ext}}(\mathfrak{B}), \mathcal{T}, G^{\text{int}}(\mathfrak{A}))$ be a classical reduction with \mathfrak{A} and \mathfrak{B} both being classical poly-time machines. Let $\hat{\mathfrak{A}} = \hat{\mathfrak{B}}$ be quantum poly-time machines. If \mathcal{R} is black-box straight-line, $\hat{\mathfrak{A}}$ -compatible and value-dominating, then \mathcal{R} is $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -closed.*

Proof. For any $\mathcal{A} \in E_{G^{\text{int}}}(\hat{\mathfrak{A}})$, let $\hat{\mathcal{A}} \in \hat{\mathfrak{A}}$ be such that \mathcal{A} and $\hat{\mathcal{A}}$ are G^{int} -equivalent. We argue that $T^{\mathcal{A}}$ and $T^{\hat{\mathcal{A}}}$ are G^{ext} -equivalent and hence $T^{\mathcal{A}} \in E_{G^{\text{ext}}}(\hat{\mathfrak{B}})$. Since \mathcal{A} and $\hat{\mathcal{A}}$ are G^{int} -equivalent and \mathcal{R} is value-dominating, $\omega_{G^{\text{ext}}}(T^{\hat{\mathcal{A}}}) = \omega_{G^{\text{ext}}}(T^{\mathcal{A}})$. $T^{\hat{\mathcal{A}}} \in \hat{\mathfrak{B}}$, i.e., it is quantum poly-time, since T is classical poly-time, and runs any oracle in straight-line. Finally, note that we need the compatibility condition so that all objects above are well-defined. \square

Combine the extendibility condition, we get the corollary below from Theorem 4.4.

Corollary 4.6. *Let \mathcal{R} be a classical black-box reduction for classical poly-time players. Let $\hat{\mathfrak{A}} = \hat{\mathfrak{B}}$ be quantum poly-time machines. If \mathcal{R} is $(\beta, \hat{\mathfrak{A}})$ -extendible, straight-line, and value-dominating, then \mathcal{R} is β - $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -respectful. As a consequence, there is a reduction $\hat{\mathcal{R}}(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ that is $(\beta, \hat{\mathfrak{A}})$ -effective.*

Note that in this scenario, $\hat{\mathcal{R}}$ is also straight-line and $\hat{\mathcal{T}}(\hat{\mathcal{A}}) = T^{\hat{\mathcal{A}}}$. Loosely speaking, the very same reduction carries over to the quantum setting.

4.2 Lifting Game-Updating Reductions

Sometimes we need to update \hat{G}^{ext} or \hat{G}^{int} or both, in order to capture the right computational assumption and the security property against quantum players. In this case, the classical transformation procedure may become totally inappropriate and give little clue about how to restore a quantum reduction (if there exists one).

We view this issue as a matter of “language-barrier”. One way to establish a reduction $\hat{\mathcal{R}}(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ is to introduce an *interpreter* $\hat{\mathcal{I}}$ that translates the “languages” between the players in the original (classical) and updated (quantum) games. Namely, $\hat{\mathcal{I}}$ translates an adversary $\hat{\mathcal{A}}$ in \hat{G}^{int} to an adversary $\hat{\mathcal{A}}'$ in the classical game G^{int} . Then we can reduce the issue to the game-preserving case and consider a class of quantum adversaries $\hat{\mathfrak{A}}' := \hat{\mathcal{I}}(\hat{\mathfrak{A}})$. Suppose we can lift the classical reduction to work with adversaries in $\hat{\mathfrak{A}}'$, then we end up with a quantum adversary in game G^{ext} . Next, by the same token, $\hat{\mathcal{I}}$ translates the adversary into a quantum one compatible in \hat{G}^{ext} . This procedure gives a quantum transformer $\hat{\mathcal{T}} := \hat{\mathcal{I}} \circ \hat{\mathcal{T}}_0 \circ \hat{\mathcal{I}}$ that operates as follows

$$\hat{\mathcal{A}} \in \hat{\mathfrak{A}} \xrightarrow{\hat{\mathcal{I}}} \hat{\mathcal{A}}' \xrightarrow{\hat{\mathcal{T}}_0} \hat{\mathcal{T}}_0(\hat{\mathcal{A}}') \xrightarrow{\hat{\mathcal{I}}} \hat{\mathcal{B}} \in \hat{\mathfrak{B}}.$$

We formalize this idea, and propose *class-translatable* reductions as a template for quantum-friendly reductions in the game-updating case. For simplicity, we assume only G^{int} is updated to \hat{G}^{int} and G^{ext} stays

the same. We want to investigate if a reduction of the form $(G^{\text{ext}}(\mathfrak{B}), \hat{\mathcal{T}}, \hat{G}^{\text{int}}(\hat{\mathfrak{A}}))$ can be derived. It is straightforward to adapt the treatment to the scenario where G^{ext} (or both) gets updated.

Definition 4.7 ((β, β') - $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -translatable reductions). Let \mathcal{R} be a classical reduction $(G^{\text{ext}}(\mathfrak{B}), \mathcal{T}, G^{\text{int}}(\mathfrak{A}))$ and β, β' be two functions. Let \hat{G}^{int} be a quantum game, and $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ be classes of quantum machines. We say \mathcal{R} is (β, β') - $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -translatable, if there exists a machine (i.e. Interpreter) $\hat{\mathcal{I}}$, such that the following hold:

- \mathcal{R} is β - $(\hat{\mathfrak{B}}, \hat{\mathfrak{A}}')$ -respectful, where $\hat{\mathfrak{A}}' := \hat{\mathcal{I}}(\hat{\mathfrak{A}})$.
- $(G^{\text{int}}, \hat{\mathcal{I}}, \hat{G}^{\text{int}})$ is a $(\beta', \hat{\mathfrak{A}})$ -effective reduction. Namely $\forall \hat{A} \in \hat{\mathfrak{A}}, \omega_{G^{\text{int}}}(\hat{\mathcal{I}}(\hat{A})) \geq \beta'(\omega_{\hat{G}^{\text{int}}}(\hat{A}))$.

Theorem 4.8 (Quantum lifting for game-updating reductions). *If $\mathcal{R}(\mathfrak{A}, \mathfrak{B})$ is (β, β') - $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -translatable, then there exists an $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -consistent reduction $\hat{\mathcal{R}}(\hat{\mathfrak{A}}, \hat{\mathfrak{B}}) := (G^{\text{ext}}(\hat{\mathfrak{B}}), \hat{\mathcal{T}}, \hat{G}^{\text{int}}(\hat{\mathfrak{A}}))$ that is $(\beta \cdot \beta', \hat{\mathfrak{A}})$ -effective.*

Proof. By the hypothesis, we know there is an interpreter $\hat{\mathcal{I}}$. Since \mathcal{R} is β - $(\hat{\mathfrak{B}}, \hat{\mathfrak{A}}')$ -respectful, by Theorem 4.4, there is a $\hat{\mathcal{T}}_0$ s.t. $(G^{\text{ext}}(\hat{\mathfrak{B}}), \hat{\mathcal{T}}_0, G^{\text{int}}(\hat{\mathfrak{A}}'))$ is $(\beta, \hat{\mathfrak{A}}')$ -effective. Define $\hat{\mathcal{T}} := \hat{\mathcal{T}}_0 \circ \hat{\mathcal{I}}$ and $\hat{\mathcal{R}} := (G^{\text{ext}}, \hat{\mathcal{T}}, \hat{G}^{\text{int}})$. Clearly, $\hat{\mathcal{R}}$ is $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -consistent because for any $\hat{A} \in \hat{\mathfrak{A}}, \hat{\mathcal{T}}(\hat{A}) = \hat{\mathcal{T}}_0(\hat{\mathcal{I}}(\hat{A})) \in \hat{\mathfrak{B}}$. On the other hand, for any $\hat{A} \in \hat{\mathfrak{A}}$ it holds that $\omega_{G^{\text{ext}}}(\hat{\mathcal{T}}(\hat{A})) \geq \beta \cdot \omega_{G^{\text{int}}}(\hat{\mathcal{I}}(\hat{A})) \geq \beta\beta' \cdot \omega_{\hat{G}^{\text{int}}}(\hat{A})$. \square

In contrast to the game-preserving setting, applying lifting theorem for game-updating reductions typically needs non-trivial extra work. The main difficulty comes from showing existence of an interpreter $\hat{\mathcal{I}}$ with the desired properties. In Sect. 5.2, we give an example that demonstrates potential applications of Theorem 4.8.

5 Applications

We give a few examples to demonstrate our framework for “quantum-friendly” reductions. In the game-preserving setting (Section 5.1), we show two versions of quantum-secure hash-based signatures schemes assuming quantum-resistant one-way functions. One follows the generic construction that builds upon Lamport’s OTS and Merkle’s original hash-tree idea. The other is an efficient variant proposed in [BDH11] that uses a more compact one-time signature scheme and a more sophisticated tree structure. In the game-updating setting (Section 5.2), we give an alternative proof for Full-Domain Hash (FDH) in the Quantum RO model as shown in [Zha12b]. We stress that this proof is meant to illustrate how our lifting theorem can be potentially applied, as apposed to providing new technical insights. Unless otherwise specified, all players are either classical or quantum poly-time machines.

5.1 Quantum Security for Hash-based Signatures

Classically, there are generic constructions (and efficient variants) for EU-CMA-secure signature schemes based on one-way functions. We show that security reductions there can be lifted easily, using our *class-respectful* characterization. It follows that there are classical signature schemes that are secure against quantum attacks, merely assuming existence of quantum-resistant one-way functions.

5.1.1 Generic hash-tree signature schemes

A generic approach for constructing EU-CMA-secure signature scheme from OWFS goes as follows:

- A one-time signature (OTS) is constructed based on OWFS. There are various ways to achieve it. We consider Lamport’s construction (L-OTS) here [Lam79].

- A family of universal one-way hash functions (UOWHFS) is constructed based on OWFS. This was shown by Rompel [Rom90] and we denote the hash family R-H.
- An OTS scheme is converted to a full-fledged (stateful) signature scheme using UOWHFS. The conversion we consider here is essentially Merkle’s original hash-tree construction [Mer90].

We show next that each step can be “lifted” to the quantum setting using our lifting theorem for game-preserving reductions (Theorem 4.4) and the straight-line characterization (Theorem 4.5). Note that we do not intend to optimize the construction here. For instance, one can use a pseudorandom function to make the signature scheme stateless. Verifying whether these still hold in the quantum setting is left as future work, though we believe it is the case, following the framework and tools we have developed.

LAMPORT’S OTS. Consider the (classical) reduction $\mathcal{R} := (G^{\text{INV}}, \mathcal{T}, G^{\text{OT-FOR}})$, where G^{INV} is the inversion game and $G^{\text{OT-FOR}}$ is the one-time forgery game. It is straight-line and value-dominating. Both games are compatible with \mathcal{Q} and $\omega_{\text{Gext}}(T^{\mathcal{A}}) \geq \beta \cdot \omega_{\text{Gint}}(\mathcal{A})$ for any \mathcal{A} with $\beta(x) = \frac{1}{2^{\ell(n)}}x$ and $\ell(n)$ a polynomial representing the length of the messages. Hence \mathcal{R} is (β, \mathcal{Q}) -effective as well. Thus we claim that:

Proposition 5.1. $(G^{\text{INV}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{OWF}}$ implies $(G^{\text{OT-FOR}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{L-OTS}}$. Namely, assuming quantum-resistant OWFS, there exists EU-CMA-secure OTS against quantum attackers \mathcal{Q} .

UOWHFS FROM OWFS. Rompel’s construction is complicated and the proof is technical (or rather tedious). However, the key ingredients in which security reductions are crucial are actually not hard to check. Basically, there are four major components in the construction:

1. From a given OWF f^0 , construct another OWF f with certain structure. Basically, f is more “balanced” in the sense that sampling a random element in the range of f and then sub-sampling its pre-images is not much different from sampling a random element in the domain directly.
2. From f , construct $\mathcal{H} = \{h_s\}$ such that for any x , it is hard to find a collision in the so called “hard-sibling” set. The hard-sibling set should comprise a noticeable fraction of all possible collisions.
3. Amplifying the hard-sibling set so that finding any collision of a pre-determined x is hard.
4. Final refinements such as making the hash functions compressing.

The second step is the crux of the entire construction. There are three reductions showing that finding a hard-sibling is as hard as inverting f which we will discuss in a bit detail below, whereas showing that the hard-sibling set is noticeably large is done by a probabilistic analysis and holds information-theoretically. Other steps either do not involve a security reduction and relies purely on some probabilistic analysis, or the reductions are clearly liftable.

The three reduction in step 2 involve four games: G^{INV} —the standard inversion game for OWFS; $G^{\text{INV}'}$ —a variant of G^{INV} in which y is sampled according to another distribution, as opposed to sampling a domain element x uniformly at random and setting $y := f(x)$; $G^{\text{COL}'}$, a variant of the collision game for UOWHFS, in which an adversary is supposed to find a collision x' in a special set (we don’t specify it here); and $G^{\text{COL}''}$, which further modifies $G^{\text{COL}'}$ in the distribution that s is sampled (instead of uniformly at random). Then $\mathcal{R}_1 = (G^{\text{INV}}, \mathcal{T}_1, G^{\text{INV}'})$, $\mathcal{R}_2 := (G^{\text{INV}'}, \mathcal{T}_2, G^{\text{COL}'})$ and $\mathcal{R}_3 = (G^{\text{COL}'}, \mathcal{T}_3, G^{\text{COL}''})$ are constructed. \mathcal{R}_1 and \mathcal{R}_3 essentially follow from the “balanced” structure of f , and \mathcal{R}_2 comes from the construction of $\mathcal{H} = \{h_s\}$. All three reductions are black-box straight-line, value-dominating, and (β_i, \mathcal{Q}) -effective with $\beta_i \geq 1/p_i(n)$ for some polynomial $p_i, i \in \{1, 2, 3\}$. For concreteness, we can set $p_1 = \ell'(n)$ —the length of the input string of f^0 , $p_2 = 3$ a constant, and $p_3(n) = 5\ell'(n) + \log \ell'(n) + 2$. Our exposition here and parameter choices are adapted from [KK05].

Proposition 5.2. $(G^{\text{INV}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{OWF}}$ implies $(G^{\text{COL}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{R-H}}$. Namely, assuming quantum-resistant OWFS, there exist UOWHFS secure against quantum attackers \mathcal{Q} .

HASH-TREE: CONVERTING OTS TO FULL-FLEDGED SIGNATURES. Once a family of UOWHFS and an OTS are at hand, we can get a full-fledged signature scheme based on Merkle’s hash-tree construction. Basically, one constructs a depth- k binary tree and each leaf corresponds to a message. Each node in the tree is associated with a key-pair (pk_w, sk_w) of the OTS scheme. The signature of a message m consists of $\sigma_m := \text{Sign}(sk_m, m)$ and an authentication chain. For each node w along the path from the root to the message, we apply $\mathcal{H} = \{h_s\}$ to the concatenation of its children’s public keys and then sign the resulting string with its secret key sk_w . The authentication chain contains all these $(pk_{w0}, pk_{w1}, \sigma_w := \text{Sign}(sk_w, pk_{w0} || pk_{w1}))$. Let M-TREE be the resulting tree-based scheme and G^{FOR} be the forgery game. The classical security analysis builds upon two reductions $(G^{\text{COL}}, \mathcal{T}, G^{\text{FOR}})$ and $(G^{\text{OT-FOR}}, \mathcal{T}', G^{\text{FOR}})$. It is easy to check that both satisfy the conditions in Corollary 4.6.

Proposition 5.3. $(G^{\text{COL}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{UOWHFS}}$ and $(G^{\text{OT-FOR}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{OTS}}$ imply $(G^{\text{FOR}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{M-TREE}}$. Namely, assuming quantum-resistant UOWHFS and OTS, there exist an EU-CMA-secure signature scheme against quantum attackers \mathcal{Q} .

Combining Propositions 5.1, 5.2, and 5.3, we get

Theorem 5.4. Assuming quantum-resistant OWFS, there exists EU-CMA-secure signature schemes against quantum poly-time attackers \mathcal{Q} .

5.1.2 XMSS: an efficient variant

The XMSS scheme [BDH11] can be seen an efficient instantiation of the generic construction above. It uses a different one-time signature scheme called Winternitz-OTS (w-OTS for short), which can be based on a family of pseudorandom functions, which in turn exists from the “minimal” assumption that OWFS exist. The hash-tree (which is called XMSS-tree in [BDH11]) also differs slightly. We now show that both the security of w-OTS and the conversion by XMSS-tree are still valid against quantum adversaries.

QUANTUM SECURITY OF W-OTS. Classically, existence of OWF imply the EU-CMA-security of w-OTS. This is established in three steps: 1) By standard constructions, a pseudorandom generator (PRG) can be constructed from OWFS [HILL99], and then one can construct a pseudo-random function (PRF) from a PRG [GGM86]. 2) A PRF is shown to be also key-one-way (KOW, defined later). 3) Show that KOW implies EU-CMA-security of w-OTS by a reduction.

The first step is known to be true in the presence of quantum adversaries [Zha12a]⁴. Informally the game for KOW of a function family F goes as follows: \mathcal{C} samples a random function $f_k \in_R F$ and a random element x in the domain. $(x, y := f_k(x))$ is sent to an adversary \mathcal{A} , who tries to find k' such that $f_{k'}(x) = y$. The PRF to KOW reduction is straight-line and value-dominating. Extendibility is trivial. Therefore it is \mathcal{Q} -respectful. This is also the case in the KOW to EU-CMA-security of w-OTS reduction. In addition β is 1 for both reductions, which means that the effectiveness (i.e., tightness in terms of success probability) in the classical analysis carries over unchanged to the quantum setting.

Proposition 5.5. $(G^{\text{PRF}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))$ implies $(G^{\text{OT-FOR}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{w-OTS}}$. Namely, assuming a quantum-resistant PRF, w-OTS is one-time EU-CMA-secure against quantum attackers \mathcal{Q} .

⁴It is easy to verify that the security reduction from PRG to PRF in GMM construction is quantum friendly. The security analysis in the HILL PRF construction from OWFS is much more complicated. To the best of our knowledge, no rigorous argument has appeared in the literature. It would be a nice exercise to apply our framework and give a formal proof.

XMSS-TREE. The XMSS-tree modifies Merkle’s hash-tree construction with an XOR-technique. Loosely speaking, each level of the tree is associated with two random strings, which mask the two children nodes before we apply the hash function to produce an authentication of a node. This tweak allows one to use a second-preimage resistant (SPR) hash function, instead of collision-resistant hash functions or UOWHFS. Theoretically universal one-way implies second-preimage resistance. But in practice people typically test second-preimage resistance when a hash function is designed. Despite this change, the security proof is not much different. Reductions are given that convert a forger either to a forger for w-OTS or to an adversary that breaks SPR-hash functions. They are straight-line, value-dominating and $(1, \mathcal{Q})$ -extendible. By Corollary 4.6, we have

Proposition 5.6. $(G^{\text{SPR}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))$ and $(G^{\text{OT-FOR}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{w-OTS}}$ imply $(G^{\text{FOR}}, \varepsilon_{\mathcal{Q}} = \text{negl}(n))_{\text{XMSS}}$. Namely, assuming quantum-resistant PRF and SPR hash functions, XMSS signature is EU-CMA-secure against quantum attackers \mathcal{Q} .

As mentioned above, UOWHFS are by definition second-preimage resistant. As a result, quantum-resistant SPR hash functions can be constructed from quantum-resistant OWFS as well. Thus, we obtain that the XMSS signature scheme is EU-CMA-secure against efficient quantum attackers \mathcal{Q} , assuming quantum-resistant OWFS.

5.2 Full-Domain Hash in Quantum Random-Oracle Model

Full domain hash (FDH) is a generic approach of constructing signature schemes based on trapdoor permutations (TDPs) in the RO model [BR93]. The classical proof cleverly “programs” the random-oracle, so that a challenge of inverting a TDP gets embedded as one output value of the random-oracle. However when we consider FDH in the quantum random-oracle (QRO) model, in which one can query the random-oracle in superposition, we lose the “programmable” ability in the proof. Zhandry [Zha12b] resolved this issue by some quantum “programming” strategy, which built upon lower bounds on quantum query complexity. This is summarized as follows.

Theorem 5.7 ([Zha12b, Theorem 5.3]). *Let F be a quantum-resistant trapdoor permutation. If we model H as a quantum random-oracle, then Π is quantum EU-CMA-secure.*

We note that Zhandry’s proof fits our framework for lifting game-updating reductions. Namely, let G^{TDP} the inversion game for a TDP. We can construct an interpreter $\hat{\mathcal{I}}$ for any adversary in the forgery game $G^{\text{QRO-FOR}}$, and show that the classical reduction $(G^{\text{TDP}}, \mathcal{T}, G^{\text{RO-FOR}})$ is translatable. Applying Theorem 4.8 proves the theorem here. We describe a proof in Appendix A for completeness. This illustrates how to apply our framework and get (in our opinion) more modular security analysis.

5.3 Quantum Security of Classical Cryptographic Protocols

So far, we have been focusing on basic cryptographic primitives such as UOWHFS and signatures. However, our framework is not limited to these scenarios, and actually can be applied to analyzing more complicated cryptographic protocols as well. Specifically an abstraction called simple-hybrid arguments, which characterize a family of classical proofs for two-party secure computation protocols in the computational setting that go through against quantum adversaries [HSS11], can be derived easily in our framework. We defer the details in Appendix B.

6 Discussions

We have proposed a general framework to study which security reductions are quantum-friendly. The lifting theorems we developed can be used to analyze security against computationally bounded quantum adver-

saries for post-quantum cryptography. As an application, we have shown the quantum security of a generic hash-tree based signature scheme and an efficient variant (which is a promising candidate for post-quantum signature schemes to be implemented in practice).

However, this note concerns mostly the feasibility of lifting classical security proofs to the quantum setting, and there are many important aspects missing and many interesting directions to be investigated. For example, we did not consider much about the “quality” of the resulting proofs for quantum adversaries. Say, can we preserve the tightness of the classical reduction when we lift it? Tightness of security reduction is of great practical impact. Not only it affects how to set the parameters in implementations, it may render security meaningless in some cases [CMS12]. Interestingly, there are also examples where we get tighter reduction in the quantum setting, as demonstrated in the quantum Goldreich-Levin theorem [AC02]. This is also a nice example of game-updating reductions beyond the QRO model. Along the same line, another game-updating reduction that is fundamental in cryptography arises from constructing a pseudorandom permutation (PRP) from a pseudorandom function (PRF). It is not clear if the classical construction remains valid if the game defining PRP allows superposition queries to distinguish it from a truly random permutation.

There are many concrete questions left for quantum-secure signature schemes as well. We showed a quantum EU-CMA-secure signature scheme based on quantum-resistant OWFS. Can we make it strongly-unforgeable? The XMSS scheme is also known to be *forward*-secure. Is it still true against quantum adversaries? We believe both answers are positive, by similar analysis from this note. Moreover, there are generic transformations that augments a signature scheme with stronger security guarantees (e.g., from EU-CMA-secure to SU-CMA-secure). Do they hold in the quantum setting? We also note that the applications we have shown in the game-updating case are not very exciting in the sense that designing an interpreter appears no easier than coming up with a quantum reduction directly. It is helpful to further explore along this line to find more interesting applications.

Finally, we remark that quantum attacks could reduce the security level of a system, using for example Grover’s quantum search algorithm. Although not covered in this note, this issue needs to be addressed with care as well.

Acknowledgments

The author is grateful to Michele Mosca for encouraging him to write up this note. F.S. would like to thank the anonymous reviewers for valuable comments and John Schank for joyful discussions on lattice-based signature schemes. F.S. acknowledges support from Ontario Research Fund, Industry Canada and CryptoWorks21.

References

- [AC02] Mark Adcock and Richard Cleve. A quantum Goldreich-Levin theorem with cryptographic applications. In *STACS 2002*, pages 323–334. Springer, 2002.
- [BBD09] Daniel J Bernstein, Johannes Buchmann, and Erik Dahmen. *Post-quantum cryptography*. Springer, 2009.
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology–ASIACRYPT 2011*, pages 41–69. Springer, 2011.

- [BDH11] Johannes Buchmann, Erik Dahmen, and Andreas Hülsing. XMSS—a practical forward secure signature scheme based on minimal security assumptions. In *Post-Quantum Cryptography*, pages 117–129. Springer, 2011.
- [BGHB11] Gilles Barthe, Benjamin Grégoire, Sylvain Héraud, and Santiago Zanella Béguelin. Computer-aided security proofs for the working cryptographer. In *Advances in Cryptology—CRYPTO 2011*, pages 71–90. Springer, 2011.
- [BGZB09] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. Formal certification of code-based cryptographic proofs. *ACM SIGPLAN Notices*, 44(1):90–101, 2009.
- [Bla08] Bruno Blanchet. A computationally sound mechanized prover for security protocols. *Dependable and Secure Computing, IEEE Transactions on*, 5(4):193–207, 2008.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM conference on Computer and communications security*, pages 62–73. ACM, 1993.
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426. Springer, 2006.
- [BZ13] Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In *Advances in Cryptology—CRYPTO 2013*, pages 361–379. Springer, 2013.
- [CMS12] Sanjit Chatterjee, Alfred Menezes, and Palash Sarkar. Another look at tightness. In *Selected Areas in Cryptography*, pages 293–319. Springer, 2012.
- [CSST11] Claude Crépeau, Louis Salvail, Jean-Raymond Simard, and Alain Tapp. Two provers in isolation. In *Advances in Cryptology—ASIACRYPT 2011*, pages 407–430. Springer, 2011.
- [DL09] Ivan Damgård and Carolin Lunemann. Quantum-secure coin-flipping and applications. In *Advances in Cryptology—ASIACRYPT 2009*, pages 52–69. Springer, 2009.
- [FKS⁺13] Serge Fehr, Jonathan Katz, Fang Song, Hong-Sheng Zhou, and Vassilis Zikas. Feasibility and completeness of cryptographic tasks in the quantum world. In *Theory of Cryptography*, pages 281–296. Springer, 2013.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.
- [GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of computer and system sciences*, 28(2):270–299, 1984.
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, 1988.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 99–108. ACM, 2011.
- [Hal05] Shai Halevi. A plausible approach to computer-aided cryptographic proofs. *Cryptology ePrint Archive, Report 2005/181*, 2005.

- [Hal07] Sean Hallgren. Polynomial-time quantum algorithms for Pell’s equation and the principal ideal problem. *J. ACM*, 54(1):1–19, 2007.
- [HH09] Iftach Haitner and Thomas Holenstein. On the (im) possibility of key dependent encryption. In *Theory of Cryptography*, pages 202–219. Springer, 2009.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [HSS11] Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In Phillip Rogaway, editor, *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 411–428. Springer, 2011.
- [KK05] Jonathan Katz and Chiu-Yuen Koo. On constructing universal one-way hash functions from arbitrary one-way functions. *IACR Cryptology ePrint Archive*, 2005:328, 2005.
- [KL07] Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography: principles and protocols*. CRC Press, 2007.
- [KR01] Joe Kilian and Phillip Rogaway. How to protect DES against exhaustive key search (an analysis of DESX). *Journal of Cryptology*, 14(1):17–35, 2001.
- [Lam79] Leslie Lamport. Constructing digital signatures from a one-way function. *Tech. Report: SRI International Computer Science Laboratory*, 1979.
- [LN11] Carolin Lunemann and Jesper Buus Nielsen. Fully simulatable quantum-secure coin-flipping and applications. In Abderrahmane Nitaj and David Pointcheval, editors, *AFRICACRYPT*, volume 6737 of *Lecture Notes in Computer Science*, pages 21–40. Springer, 2011.
- [Lyu08] Vadim Lyubashevsky. Lattice-based identification schemes secure under active attacks. In *Public Key Cryptography–PKC 2008*, pages 162–179. Springer, 2008.
- [Lyu09] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *Advances in Cryptology–ASIACRYPT 2009*, pages 598–616. Springer, 2009.
- [Mer90] Ralph C Merkle. A certified digital signature. In *Advances in Cryptology–CRYPTO 1989*, pages 218–238. Springer, 1990.
- [MR09] Daniele Micciancio and Oded Regev. Lattice-based cryptography. In *Post-quantum cryptography*, pages 147–191. Springer, 2009.
- [Nao03] Moni Naor. On cryptographic assumptions and challenges. In *Advances in Cryptology–CRYPTO 2003*, pages 96–109. Springer, 2003.
- [Pas11] Rafael Pass. Limits of provable security from standard assumptions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 109–118. ACM, 2011.
- [Pei09] Chris Peikert. Some recent progress in lattice-based cryptography. In *TCC*, volume 5444, page 72, 2009.
- [Rom90] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 387–394. ACM, 1990.

- [Sen11] Nicolas Sendrier. Code-based cryptography. In *Encyclopedia of Cryptography and Security*, pages 215–216. Springer, 2011.
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.
- [Sho05] Victor Shoup. Sequences of games: a tool for taming complexity in security proofs. *Cryptology ePrint Archive, Report 2004/332*, 2005.
- [Stu09] Aaron Stump. Proof checking technology for satisfiability modulo theories. *Electronic Notes in Theoretical Computer Science*, 228:121–133, 2009.
- [Unr10] Dominique Unruh. Universally composable quantum multi-party computation. In *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 486–505. Springer, May 2010. Preprint on arXiv:0910.2912 [quant-ph].
- [Unr12] Dominique Unruh. Quantum proofs of knowledge. In *Eurocrypt 2012*, volume 7237 of *LNCS*, pages 135–152. Springer, April 2012. Preprint on IACR ePrint 2010/212.
- [vdG97] Jeroen van de Graaf. Towards a formal definition of security for quantum protocols. *PhD thesis, Département d’informatique et de recherche opérationnelle, Université de Montréal*, 1997.
- [Wat09] John Watrous. Zero-knowledge against quantum attacks. *SIAM J. Comput.*, 39(1):25–58, 2009. Preliminary version in STOC 2006.
- [Yao82] Andrew C Yao. Theory and application of trapdoor functions. In *Foundations of Computer Science, 1982. SFCS’08. 23rd Annual Symposium on*, pages 80–91. IEEE, 1982.
- [Zha12a] Mark Zhandry. How to construct quantum random functions. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 679–687. IEEE, 2012.
- [Zha12b] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In *Advances in Cryptology–CRYPTO 2012*, pages 758–775. Springer, 2012.

A (Alternative) Proof of Theorem 5.7: FDH in QRO

We first review a technical tool in [Zha12b] called *semi-constant* distribution. Loosely speaking, it allows us to “program” a function, which still looks like a random function even to a quantum observer.

Definition A.1 (Semi-Constant Distribution [Zha12b, Definition 4.1]). Let X and Y be sets and denote $\mathcal{H}_{X,Y}$ the set of functions from X to Y . The semi-constant distribution SC_λ is defined as the distribution over $\mathcal{H}_{X,Y}$ resulting from the following process:

- Pick a random element y from Y .
- For each $x \in X$, set $H(x) = y$ with probability λ . Otherwise set $H(x)$ to be a random element in Y .

Theorem A.2 ([Zha12b, Corollary 4.3]). *The distribution of the output of a quantum algorithm making q_H queries to an oracle drawn from SC_λ is at most a distance $\frac{8}{3}q_H^4\lambda^2$ away from the case when the oracle is drawn uniformly from $\mathcal{H}_{X,Y}$.*

We are now ready to give a proof for Theorem 5.7 using our framework for game-updating reductions.

Proof. Classically there is $\mathcal{R} = (G^{\text{TDP}}, \mathcal{T}, G^{\text{RO-FOR}})$ that inverts the TDP with a forger for the FDH-Sign scheme. We construct an interpreter $\hat{\mathcal{I}}$ as follows (Fig. 1), and show that \mathcal{R} is $(\hat{\mathcal{A}}, \hat{\mathcal{B}})$ -translatable with $\hat{\mathcal{A}} = \hat{\mathcal{B}} = \mathcal{Q}$.

Clearly, $\hat{\mathcal{A}}$ is a well-defined (quantum) adversary for the original forgery game $G^{\text{RO-FOR}}$ (i.e., the hash queries are classical). If $\hat{\mathcal{A}}$ outputs a valid forgery (m^*, σ^*) such that $\hat{H}(m^*) = f_{pk}(\sigma^*)$ and $\mathcal{O}_2(m^*) = 1$, we know that $\hat{H}(m^*) = b = H(a)$ and hence (a, σ^*) forms a valid forgery in the classical forgery game. Note that the view of $\hat{\mathcal{A}}$ in $\hat{\mathcal{A}}$ differs from a true interaction with a challenger in game $G^{\text{QRO-FOR}}$ in two places: a truly random oracle is replaced by \hat{H} drawn from SC_λ and the signing query fails with probability λ . By picking λ a proper inverse polynomial in q_H and q_S , we can obtain from Theorem A.2 that $\omega_{G^{\text{RO-FOR}}}(\hat{\mathcal{I}}(\hat{\mathcal{A}})) \geq \omega_{G^{\text{QRO-FOR}}}^2(\hat{\mathcal{A}})/p(n)$ for some polynomial $p(\cdot)$. Thus $(G^{\text{RO-FOR}}, \hat{\mathcal{I}}, G^{\text{QRO-FOR}})$ forms a (β', \mathcal{Q}) -effective reduction for a suitable β' . Since the two random oracles $(\mathcal{O}_1, \mathcal{O}_2)$ can be simulated efficiently by k -wise indecent functions (C.f. [Zha12b, Theorem 6.1]), \mathcal{R} is clearly β - $(\mathcal{Q}, \hat{\mathcal{I}}(\mathcal{Q}))$ -respectful with $\beta = 1$. Therefore we obtain that \mathcal{R} is $(\mathcal{Q}, \mathcal{Q})$ -translatable, which by Theorem 4.8 can be lifted to a reduction $(G^{\text{TDP}}(\mathcal{Q}), \hat{\mathcal{T}}, G^{\text{QRO-FOR}}(\mathcal{Q}))$. This shows that the FDH-Signature scheme is quantum EU-CMA-secure, assuming quantum-resistant trapdoor permutations.

Interpreter $\hat{\mathcal{I}}$

Input: Adversary $\hat{\mathcal{A}}$ for a quantum EU-CMA-game. Let q_S and a_H be upper bounds on the number of signing queries and hash queries of $\hat{\mathcal{A}}$.

Output: An adversary $\hat{\mathcal{A}}' := \hat{\mathcal{I}}(\hat{\mathcal{A}})$ that operates as follows:

1. Receive pk from a challenger, which indexes a permutation f_{pk} .
2. Pick an arbitrary message a . Query $H(\cdot)$ and get $b := H(a)$.
3. Emulate (internally) a quantum EU-CMA-game with $\hat{\mathcal{A}}$.
 - Use b to create an oracle \hat{H} from a semi-constant distribution SC_λ which handles (quantum) hash queries from $\hat{\mathcal{A}}$. Specifically, let \mathcal{O}_2 be a random oracle outputting 1 with probability λ and \mathcal{O}_1 be a random oracle mapping a message to an input of f_{pk} . Let $\hat{H}(x) = b$ if $\mathcal{O}_2(x) = 1$ and $\hat{H}(x) = f_{pk}(\mathcal{O}_1(x))$ otherwise.
 - On signing query m_i , if $\mathcal{O}_2(m_i) = 1$ abort. Otherwise respond with $\sigma_i := \mathcal{O}_1(m_i)$.
4. On output (m^*, σ^*) from $\hat{\mathcal{A}}$, if $\mathcal{O}_2(m^*) = 1$ output (a, σ^*) .

Figure 1: Construction of the Interpreter. □

B Details on Sect. 5.3

Security definitions in this setting usually follows the *simulation paradigm*. In particular, there is not a simple game capturing them⁵. Roughly speaking, we require the existence of an imaginary entity (called the simulator) with certain properties for any possible adversary. The main ingredient of a security proof is often a hybrid argument, in which a sequence of imaginary experiments (a.k.a. *hybrids*) are defined in terms of an adversary and the simulator. The goal is to show each adjacent pair of hybrids is indistinguishable. Whenever this is done by a reduction of breaking a computational assumption, we can define a distinguishing game (as our internal game) and study if the reduction can be lifted using our framework.

⁵In some sense, the security definitions we discussed earlier that are specified by games are *falsifiable*, which does not seem to be so here.

Consider zero-knowledge proof protocols as a concrete example. Zero-knowledge property requires that for any dishonest verifier V^* , there is a simulator \mathcal{S} , such that the output of \mathcal{S} is indistinguishable from the view of V^* in real protocol with honest prover. At this moment, it looks quite alien to our framework. However, once we start the security proof, it naturally fits our framework. Basically, if we fix a dishonest V^* , and a specific construction of a simulator, showing that the simulator works can be thought of as a distinguishing game.

ZK Distinguishing Game $G_{V^*, \mathcal{S}}^{\text{ZK}}$

Two parties: Challenger \mathcal{C} and distinguisher \mathcal{D} .

- \mathcal{C} flips a random coin $b \in_R \{0, 1\}$. If $b = 0$ simulates an execution of the ZK protocol and sends \mathcal{D} the view of V^* . If $b = 1$, run the simulator \mathcal{S} and sends \mathcal{D} the output of \mathcal{S} .
- \mathcal{D} receives the message from \mathcal{C} , generate one bit b' and send it to \mathcal{C} .
- \mathcal{C} outputs succ if $b = b'$ and fail otherwise.

The security proof will then proceed in the familiar fashion. Namely a reduction $(G^{\text{ext}}, \mathcal{T}, G^{\text{int}} := G_{V^*, \mathcal{S}}^{\text{ZK}})$ is constructed for some computational assumption captured by G^{ext} . We can then ask if we can “lift” the reduction to the quantum setting. One subtlety, however, is that the distinguishing game is specific to V^* and \mathcal{S} . Because of issues like rewinding, we have to update the games. The challenge then lies in constructing a simulator $\hat{\mathcal{S}}$ for any dishonest quantum verifier \hat{V}^* , which gives the updated distinguishing game $\hat{G}_{\hat{V}^*, \hat{\mathcal{S}}}^{\text{ZK}}$ in the presence of quantum verifiers.

Sometimes we end up in the simpler game-preserving case. A concrete example is an abstraction proposed in [HSS11], called *simple-hybrid arguments* (SHA).

SIMPLE HYBRID ARGUMENTS. SHA formalizes a family of classical proofs that can go through against quantum adversaries in the computational UC model. The essence is a simple observation: if two adjacent hybrids only differs by a small change such as chaining the plaintext of an encryption, then quantum security immediately follows as long as computational assumptions are made quantum-resistant. Using our framework, each adjacent pair of hybrid induce a distinguishing game G^{int} that can be defined similarly to $G_{V^*, \mathcal{S}}^{\text{ZK}}$, and a classical reduction $\mathcal{R} := (G^{\text{ext}}, \mathcal{T}, G^{\text{int}})$ is already at hand for some computational assumption defined by G^{ext} . The conditions in SHA, e.g., changing only the plaintext, ensure that \mathcal{R} satisfy the definition of $(\hat{\mathfrak{A}}, \hat{\mathfrak{B}})$ -respectful reductions with $\hat{\mathfrak{A}} = \hat{\mathfrak{B}} = \mathcal{Q}$. As a result, these reductions can be lifted by Theorem 4.4.